ALBERTO MARIA METELLI
Curriculum Vitae et Studiorum

Dipartimento di Elettronica, Informazione e Bioingegneria
Politecnico di Milano
32, Piazza Leonardo da Vinci
20133, Milano, Italy

✉ albertomaria.metelli@polimi.it
📞 (+39) 02 2399 9679
🌐 albertometelli.github.io

SUMMARY

Alberto Maria Metelli is an Assistant Professor of Information Processing Systems with the Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB) in the Artificial Intelligence and Robotics Laboratory (AIRLab) at Politecnico di Milano (PoliMi). He obtained the Ph.D. in Information Technology (cum Laude) in March 2021 at Politecnico di Milano defending a thesis about environment configuration in reinforcement learning, awarded the "Premio NeoDottori di Ricerca Marco Cadoli 2021" as the best Italian Ph.D. thesis in Artificial Intelligence (from AIxIA) and recently published in the book series "Frontiers in Artificial Intelligence and Applications" (FAIA) [B1]. He is co-founder of ML Cube S.r.l., an innovative start-up, providing cutting-edge solutions for machine learning systems and life-cycle-management optimization, nominated by Fortune Italia as one of the top 20 Italian AI start-ups in 2022. His main research interests revolve around artificial intelligence and machine learning for *sequential decision-making*, in particular *reinforcement learning* (RL). He is currently working on theoretical and algorithmic aspects of inverse RL [C22, C19, C4], RL in configurable environments [C29, T2, C18], off-policy RL [C28, J9, C17], automated RL [C13, J4], and RL in structured environments [C6]. He participates in research projects about reinforcement learning for autonomous driving [J10], defense, Industry 4.0, complex networks, and about machine learning for climate science [A4, A3]. He is also interested in algorithms, optimization, statistics, probability, and recommendation systems [W23].

HIGHLIGHTS

- **Assistant professor** at DEIB, Politecnico di Milano, since March 2023.
- **National scientific qualification** as Italian associate professor (sector 09/H1), until February 2034.
- Author/Co-author of 11 publications in peer-reviewed international journals (4 as main contributor, 1 single-author), including JMLR (2), Machine Learning (2), RAS (1), ESWA (1), IEEE TNNLS (1), and IEEE T-ITS (1). According to Scimago Journal Rank: **8 articles in Q1 journals**. Author/Co-author of 30 publications in peer-reviewed international conferences (11 as main contributor), including ICML (10), NeurIPS (8), AAAI (6), AISTATS (2), UAI (1). According to GGS Conference Rating: **24 publications in A++ venues** and 2 publications in A+ venues (Class 1).
- Winner of both "**Premio NeoDottori di Ricerca** Marco Cadoli 2021" and "Premio NeoLaureati Leonardo Lesmo 2018" for the best Italian Ph.D. thesis and M.Sc. thesis in Artificial Intelligence respectively, awarded by Associazione Italiana per l'Intelligenza Artificiale (AIxIA).
- **Oral presentation at NeurIPS 2018** (30/4856 submissions, 0.62%), spotlight presentation at NeurIPS 2021 (260/9122 submissions, 2.9%), notable paper at AISTATS 2023 (32/1689 submissions, 1.9%), oral presentation at ICML 2023 (156/6538 submissions, 2.39%).
- **Lecturer** of the Ph.D. course Reinforcement Learning (since a.y. 2021-2022) and of the B.Sc. course Computer Science (since a.y. 2023-2024) at Politecnico di Milano. Teaching assistant of M.Sc. course Machine Learning (since a.y. 2021-2022), B.Sc. and M.Sc. course Foundations of Artificial Intelligence (since a.y. 2021-2022) at Politecnico di Milano.
- **Associate Editor** of IEEE Robotics and Automation Letters (RA-L). **Program chair and co-organizer** of the 15th European Workshop of Reinforcement Learning (EWRL 2022). Session chair of ECML-PKDD 2021. **Senior PC member** of two top-tier conferences (AAAI 2022 and IJCAI 2021).
- Participation in 7 industrial research projects funded by private companies (1 as **co-PI**) and in 4 competitive research project funded by public institutions (1 as **co-PI** of the local research unit).
- (Co-)Supervisor of 10+ Ph.D. students and of 30+ M.Sc. students at Politecnico di Milano.
- Member of the ELLIS Society and CLAIRE supporter.
- **Co-founder** of ML Cube S.r.l., accredited spin-off of the Politecnico di Milano, born in November 2020.

Personal Information

Address:	Dipartimento di Elettronica, Informazione e Bioingegneria - Politecnico di Milano
	32, Piazza Leonardo da Vinci - 20133, Milano, Italy
Email:	albertomaria.metelli@polimi.it
Voice:	(+39) 02 2399 9679
Institutional page:	https://www.deib.polimi.it/eng/people/details/926910
Webpage:	https://albertometelli.github.io
dblp:	https://dblp.org/pid/209/4941.html
Google Scholar:	https://scholar.google.com/citations?user=R31IsPwAAAAJ
ORCID:	https://orcid.org/0000-0002-3424-5212
Scopus:	https://www.scopus.com/authid/detail.uri?authorId=57195947711 (57195947711)
Web of Science:	https://www.webofscience.com/wos/author/record/2074428 (AAY-5206-2020)

Academic Experience

Assistant Professor (Ricercatore a tempo determinato (Junior) - L.240/2010, art.24, c.3, lett.a)) *March 2023-*
Politecnico di Milano - Dipartimento di Elettronica, Informazione e Bioingegneria *Milan, Italy*
Research programme: "Artificial intelligence foundations for sequential decision making".
SSD: ING-INF/05, Settore concorsuale: 09/H1 - Sistemi di elaborazione delle informazioni.

Postdoctoral Researcher (Assegnista di ricerca - L.240/2010, art.22) *November 2020-February 2023*
Politecnico di Milano - Dipartimento di Elettronica, Informazione e Bioingegneria *Milan, Italy*
Research title: "Development of reinforcement learning algorithms for autonomous driving applications".
Research manager: Prof. Marcello Restelli.
SSD: ING-INF/05, Area: 09.

National Scientific Qualifications

National Scientific Qualification as Italian Associate Professor *February 2023*
(Abilitazione Scientifica Nazionale - L.240/2010, art.16)
Professore di II fascia - Settore concorsuale: 09/H1 - Sistemi di elaborazione delle informationi.
Validità: 06/02/2023 - 06/02/2034

National Professional Qualifications

Professional Qualification as Italian Information Engineer *September 2021*
(Abilitazione all'esercizio della professione di Ingegnere dell'Informazione - DPR.328/2001)
Sezione A - Settore dell'informazione. I Sessione 2021 - Politecnico di Milano.

Education

Ph.D. in Information Technology (Dottorato di ricerca) *November 2017-October 2020*
Politecnico di Milano - Dipartimento di Elettronica, Informazione e Bioingegneria *Milan, Italy*
Supervisor: Prof. Marcello Restelli. Tutor: Prof. Nicola Gatti.
Ph.D. Thesis: "Exploiting environment configurability in reinforcement learning".
Reviewers: Prof. Amir-massoud Farahmand (Vector Institute, University of Toronto), Prof. Alessandro Lazaric (Facebook AI, Paris).
Date of award: 11 March 2021. Final Mark: **Laude**.

M.Sc. in Computer Science and Engineering (LM-32 Ingegneria Informatica) *October 2015-July 2017*
Politecnico di Milano *Milan, Italy*
M.Sc. Thesis: "Compatible reward inverse reinforcement learning".
Supervisor: Prof. Marcello Restelli. Co-supervisor: Dott. Matteo Pirotta.
Date of award: 27 July 2017. GPA: 30/30. Final Mark: **110/110 cum Laude**.

B.Sc. in Engineering of Computing Systems (L-8 Ingegneria dell'informazione) *October 2012-July 2015*
Politecnico di Milano *Milan, Italy*
Date of award: 24 July 2015. GPA: 30/30. Final Mark: **110/110 cum Laude**.

High School Diploma *September 2007-July 2012*
IISS Ettore Majorana *Seriate (Bg), Italy*
Specialization in Computer Science (Perito Informatico). Final Mark: 100/100 cum Laude.

Publications

Scientific Productivity   **41** publications (**41** entries on Scopus, **51** co-authors according to Scopus)

- <u>International Journals</u>: Author/Co-author of **11 publications** in peer-reviewed international journals (4 as main contributor, 1 single-author), including **Journal of Machine Learning Research** (**2**), **Machine Learning** (**2**), **Robotics and Autonomous Systems** (**1**), **Expert Systems with Applications** (**1**), **IEEE Transactions on Neural Networks and Learning Systems** (**1**), and **IEEE Transactions on Intelligent Transportation Systems** (**1**).
    - According to Scimago Journal Rank (SJR):[1] **8 publications in Q1** journals (of which 6 in "Artificial Intelligence" area and 2 in "Computer Science Applications" area).
    - According to CORE Journal Ranks:[2] **3 publications in A*** journals and 2 publications in A journals.
- <u>International Conferences</u>: Author/Co-author of **30 publications** in peer-reviewed international conferences (11 as main contributor), including **ICML** (**10**), **NeurIPS** (**8**), **AAAI** (**6**), **AISTATS** (**2**), and **UAI** (**1**).
    - According to GII-GRIN-SCIE (GGS) Conference Rating:[3] **24 publications in A++** venues and 2 publications in A+ venues (Class 1).
    - According to CORE Conference Ranks:[4] **24 papers in A*** venues and 3 papers in A venues.

Publication Impact	<u>Google Scholar:</u>	citations **708**	*h*-index **15**	*i*10-index **19**
(accessed on 5 December 2023)	<u>Scopus:</u>	citations **332**	*h*-index **11**	*i*10-index **13**

Fellowships, Awards, and Recognitions

Personal Research Awards and Recognitions

AAAI 2024 - New Faculty Highlights   *December 2023*
Association for the Advancement of Artificial Intelligence (AAAI)
Description: Invited speaker program at AAAI conference (GGS: A++) highlighting AI researchers who have just begun careers as new faculty members or the equivalent in industry.
(link: https://aaai.org/aaai-conference/aaai-24-new-faculty-highlights/)

DAAD AInet Fellowship - Postdoc-NeT-AI on Human-centered AI   *October 2023*
German Academic Exchange Service (DAAD)
Description: Fellowship granted to excellent early-career researchers to meet the German AI research community.
(link: https://www.daad.de/en/the-daad/postdocnet/fellows/fellows/)

Publication of the Ph.D. Thesis as Book   *June 2022*
European Association for Artificial Intelligence (EurAI)
Description: Publication of the Ph.D. thesis in the "Dissertation in AI" series "Frontiers in Artificial Intelligence and Applications" as a recognition for the outstanding theses participating in the EurAI "Dissertation Award".
(link: https://www.iospress.com/catalog/books/exploiting-environment-configurability-in-reinforcement-learning)

Winner of "Premio NeoDottori di Ricerca Marco Cadoli 2021"   *November 2021*
Associazione Italiana per l'Intelligenza Artificiale (AIxIA)
Description: Award for the best Italian Ph.D. thesis in Artificial Intelligence (public selection). Amount: 1000€.
(link: https://aixia.it/premi/premio-per-neodottori-di-ricerca-marco-cadoli-annuale/)

Recipient of a "Springer Award"   *April 2021*
IT PhD Board of Professors, Politecnico di Milano
Description: Publication in a Polimi SpringerBriefs volume, for the best results from the IT PhD program doctors.
(link: https://link.springer.com/book/10.1007/978-3-030-85918-3)

Winner of "Premio NeoLaureati Leonardo Lesmo 2018"   *October 2018*
Associazione Italiana per l'Intelligenza Artificiale (AIxIA)
Description: Award for the best Italian M.Sc. thesis in Artificial Intelligence (public selection). Amount: 500€.
(link: https://aixia.it/premi/premio-per-neolaureati-leonardo-lesmo-annuale/)

---

[1] https://www.scimagojr.com
[2] http://portal.core.edu.au/jnl-ranks
[3] https://scie.lcc.uma.es/
[4] http://portal.core.edu.au/conf-ranks

Conference Awards and Recognitions

Oral Presentation at ICML 2023 (top 2.39%) *July 2023*
International Conference on Machine Learning 2023, Honolulu, HI, US
Description: Best 156 papers out of 6538 submissions (top 2.39%) to ICML 2023 (GGS: **A++**) for the paper
"Towards Theoretical Understanding of Reinforcement Learning" [C4].

Notable Paper at AISTATS 2023 (top 1.9%) *April 2023*
International Conference on Artificial Intelligence and Statistics 2023, Valencia, Spain
Description: Best 32 papers out of 1689 submissions (top 1.9%) to AISTAS 2023 (GGS: **A+**) for the paper "A Tale
of Sampling and Estimation in Discounted Reinforcement Learning" [C7].

Spotlight Presentation at NeurIPS 2021 (top 2.9%) *December 2021*
Neural Information Processing Systems 2021, Online
Description: Best 260 papers out of 9122 submissions (top 2.9%) to NeurIPS 2021 (GGS: **A++**) for the paper
"Subgaussian and Differentiable Importance Sampling for Off-Policy Evaluation and Learning" [C17].

Oral Presentation at NeurIPS 2018 (top 0.62%) *December 2018*
Neural Information Processing Systems 2018, Montreal, Canada
Description: Best 30 papers out of 4856 submissions (top 0.62%) to NeurIPS 2018 (GGS: **A++**) for the paper
"Policy Optimization via Importance Sampling" [C28].

Student Awards and Recognitions

Winner of a Ph.D. Scholarship *July 2017*
Ministero dell'Istruzione, dell'Università e della Ricerca (MIUR)
Description: Ranked first among 50 applicants of Politecnico di Milano in Computer Science and Engineering area.

Recipient of a "Le migliori matricole dell'A.A. 2012/2013" award *May 2014*
Politecnico di Milano
Description: Awarded for the best freshmen students of Politecnico di Milano.

B.Sc. and M.Sc. Student Scholarships *2012-2017*
Politecnico di Milano
Description: Full reduction of student tuition fees for merit GPA $> 29/30$ (each eligible year).

Enrollment to the "Albo Nazionale delle Eccellenze" *2012*
Istituto Nazionale di Documentazione Innovazione e Ricerca Educativa (INDIRE)
Description: National register of the best among the Italian high-school graduates.
(link: https://www.indire.it/eccellenze/)

Competitions

Runner-Up (Best Academic Team) at ACM RecSys Challenge 2017 *July 2017*
Association for Computing Machinery (ACM)
Description: Recommender systems challenge with a team of 8 students from Politecnico di Milano. Amount: 1500€.
(link: http://www.recsyschallenge.com/2017/)

Winner of "Migliori Elaborati 2011-2012" *June 2012*
IISS Ettore Majorana
Description: Best project award for the developing of a database system to manage remedial courses.

Winner of a Bronze medal at "Olimpiadi Italiane di Informatica 2011" *October 2011*
Associazione Italiana per l'Informatica ed il Calcolo Automatico (AICA)
Description: Italian Olympiad in Informatics, ranked 24th.
(link: https://www.olimpiadi-informatica.it/index.php/olimpiadi-italiane-2010-2011.html)

Reviewing Service Recognitions

He has been recognized for the reviewing service for international conferences: NeurIPS 2019 (top 50%), ICML 2020 (top 33%), AAAI 2021 (top 25%), ICLR 2021 (outstanding reviewer), AISTATS 2022 (top 10%), NeurIPS 2023 (top reviewer).

International Conference Grants

He has been awarded travel/accommodation/registration support for international conference participation, granted by conference organizers via conference sponsor: NIPS 2017 (Long Beach, CA, US), ICML 2018 (Stockholm, Sweden), NeurIPS 2018 (Montreal, Canada), ICML 2019 (Long Beach, CA, US), NeurIPS 2019 (Vancouver, Canada), ICML 2022 (Baltimore, MD, US), NeurIPS 2023 (New Orleans, LA, US).

Teaching Activity[5]

Primary Responsibility Teaching Activities

Lecturer and coordinator of Computer Science *2023-2024*
Politecnico di Milano *Milan, Italy*
6 CFU - B.Sc. in Civil Engineering - Campus Leonardo - II semester. Avg no. of students: ~75.
a.y. 2023-2024 (36 hours).

Lecturer and co-coordinator of Reinforcement Learning *2021-2022, 2023-2024*
Politecnico di Milano *Milan, Italy*
5 CFU - Ph.D. in Information Technology - Campus Leonardo. Avg no. of students: ~80. Co-coordinator: Prof. Marcello Restelli.
a.y. 2021-2022 (15 hours, corresponding to 3 CFU), a.y. 2023-2024 (15 hours, corresponding to 3 CFU).

Lecturer of Deep Reinforcement Learning *2021-2022*
GSSI - Gran Sasso Science Institute *Online*
Ph.D. in Computer Science. Avg no. of students: ~10.
a.y. 2021-2022 (2 hours).

Teaching Assistance Activities

Teaching Assistant of Machine Learning *2021-2024*
Politecnico di Milano *Milan, Italy*
5 CFU - M.Sc. Computer Science and Engineering and M.Sc. Mathematical Engineering - Campus Leonardo - II semester. Avg no. of students: ~100. Lecturer: Prof. Daniele Loiacono.
a.y. 2021-2022 (10 hours), a.y. 2022-2023 (20 hours), a.y. 2023-2024 (20 hours).

Teaching Assistant of Foundations of Artificial Intelligence *2021-2024*
Politecnico di Milano *Milan, Italy*
5 CFU - B.Sc. and M.Sc. Computer Science and Engineering - Campus Leonardo - I semester. Avg no. of students: ~80+80. Lecturers: Proff. Francesco Amigoni and Pier Luca Lanzi.
a.y. 2021-2022 (10+10 hours), a.y. 2022-2023 (16+16 hours), a.y. 2023-2024 (8+8 hours).

Teaching Assistant of Computer Science (Informatica) *2018-2023*
Politecnico di Milano *Milan, Italy*
8 CFU - B.Sc. in Environmental and Land Planning Engineering - Campus Leonardo - I semester. Average no. of students: ~150. Lecturer: Prof. Andrea Bonarini.
a.y. 2018-2019 (24 hours), a.y. 2019-2020 (24 hours), a.y. 2020-2021 (30 hours), a.y. 2021-2022 (30 hours), a.y. 2022-2023 (24 hours).

Laboratory Teaching Assistant of Computer Science (Informatica) *2018-2021*
Politecnico di Milano *Milan, Italy*
8 CFU - B.Sc. in Environmental and Land Planning Engineering - Campus Leonardo - I semester. Avg no. of students: ~150. Lectuter: Prof. Andrea Bonarini.
a.y. 2018-2019 (20 hours), a.y. 2019-2020 (20 hours), a.y. 2020-2021 (20 hours).

Laboratory Teaching Assistant of Computer Science A (Informatica A) *2017-2018*
Politecnico di Milano *Milan, Italy*
10 CFU - B.Sc. in Management Engineering - Campus Bovisa - I semester. Avg no. of students: ~80. Lecturer: Prof. Florian Daniel.
a.y. 2017-2018 (36 hours).

Academic Tutor *2020*
Cefriel S.c.a.r.l. *Milan, Italy*
Supervision of a student for the project of the Master CEFRIEL on AI&ML (24 hours).

---

[5]Total 68 hours of teaching and 370 hours of teaching assistance.

Evaluation Committees

Examination Committees

President of the Ph.D. Final Examination Committee *April 2023*
Ph.D. in Information Technology - Politecnico di Milano *Milan, Italy*
Candidate: Dr. Giulia Romano. Thesis: "Pricing and advertising strategies in e-commerce scenarios"
PhD Committee: Proff. Alberto Maria Metelli, Ioannis Caragiannis (Aarhus University, Denmark), Negin Golrezaei (MIT Sloan School of Management, Cambridge, Massachusetts, US).

Research Projects

Competitive Research Projects

AI4REALNET (AI for REAL-world NETwork operation) *October 2023-*
HORIZON EUROPE
Role: **Co-principal investigator** of the local research unit and **WP co-leader** (with Prof. Marcello Restelli).
Topic: Reinforcement learning for operating complex real-world network infrastructures.
Project coordinator: INESC TEC, Porto, Portugal. Funding institution: European Union. Budget: 4M€ (456k€ to Politecnico di Milano). Duration: 42 months.

iBeCHANGE *December 2023-*
HORIZON RIA
Role: Research scientist.
Topic: Develop the iBeChange platform for personalizing cancer prevention.
Project coordinator: Istituto Europeo Di Oncologia Srl (IEO), Milan, Italy. Local co-principal investigators: Proff. Francesco Trovò, Emilia Ambrosini. Funding institution: European Union. Budget: 5.9M€ (644k€ to Politecnico di Milano). Duration: 60 months.

FAIR (Future Artificial Intelligence Research) - Spoke 4: Adaptive AI *March 2023-*
Extended Partnership - National Recovery and Resilience Plan (PNRR)
Role: Research scientist.
Topic: Artificial intelligence: foundational aspects - Adaptive AI.
Spoke coordinator: Politecnico di Milano (Prof. Nicola Gatti). Funding institution: Italian Ministry of University and Research (MUR). Budget: 12M€ (6M€ to Politecnico di Milano). Duration: 36 months.

CLINT (CLImate INTelligence) *July 2021-*
HORIZON 2020
Role: Research scientist.
Topic: Extreme events detection, attribution and adaptation using machine learning.
Project coordinator: Politecnico di Milano (Prof. Andrea Castelletti). Funding institution: European Union. Budget: 6M€ (1.1M€ to Politecnico di Milano). Duration: 48 months.

Competitive Research Projects for Computational Time

PGCLaSI - Policy Gradients for Control of Large-Scale Industrial Plants *October 2023-July 2024*
CINECA Supercomputing Centre - Class C Project
Role: Principal Investigator.
Topic: Development of efficient policy gradient algorithms for large-scale industrial plants.
Grant: 80k GPU hours.

Industrial-funded Research Projects

Nuovo Pignone Tecnologie S.r.l. (Baker Hughes group) *April 2023-*
Development of reinforcement learning algorithms for the control of industrial compressors
Role: **Co-principal investigator** (with Prof. Marcello Restelli).
Budget: 62k€.

ML cube S.r.l. *December 2020-December 2022*
Life-Cycle-Management and Optimization of Machine learning algorithms in real-time biddings
Role: Research scientist. Co-principal investigators: Proff. Marcello Restelli, Nicola Gatti, Francesco Trovò. Budget: 84k€.

Nuovo Pignone Tecnologie S.r.l. (Baker Hughes group) *October 2021-June 2022*
PID controller tuning using Reinforcement Learning
Role: Research scientist.

Principal investigator: Prof. Marcello Restelli. Budget: 40k€.

Leonardo S.p.A. *May 2021-*
Machine Learning per l'Autonomia dei Velivoli
Role: Research scientist.
Co-principal investigators: Proff. Marcello Restelli and Nicola Gatti. Budget: 250k€.

Ferrari S.p.A. *December 2020-January 2021*
Reinforcement Learning Techniques for Developing Artificial Test Drivers on a F1 Simulator
Role: Research scientist.
Principal Investigator: Prof. Marcello Restelli. Budget: 80k€.

Magneti Marelli S.p.A. *March 2019-February 2020*
Decision Making based on Reinforcement Learning for Automated Driving (2nd project)
Role: Research scientist.
Principal investigator: Prof. Marcello Restelli. Budget: 60 k€. Reference Publications: [W20, J10].

Magneti Marelli S.p.A. *March 2018-February 2019*
Decision Making based on Reinforcement Learning for Automated Driving (1st project)
Role: Research scientist.
Principal investigator: Prof. Marcello Restelli. Budget: 56k€.

Student Supervision

He is currently **supervising 2 Ph.D. students**:

1. **Filippo Lazzati**, "Inverse Reinforcement Learning", Ph.D. in Information Technology (Cycle: XXXIX), DEIB, Politecnico di Milano.

2. **Alessandro Montenegro**, "Policy Gradient Methods for Industrial Applications", Ph.D. in Information Technology (Cycle: XXXIX), DEIB, Politecnico di Milano. Co-supervisor: Matteo Papini.

He is **co-supervising 11 Ph.D. students** mostly of the Information Technology programme at Politecnico di Milano.

He **supervised** (relatore) **3 M.Sc. theses** and **co-supervised** (correlatore) **33 concluded M.Sc. theses** mostly in Computer Science and Engineering at Politecnico di Milano and he is currently **(co-)supervising 10 M.Sc. students** at Politecnico di Milano, including **4** students admitted to the **Honours Programme of Politecnico di Milano** and **1** student awarded with the **"Premio NeoLaureati Leonardo Lesmo 2019"** (AIxIA) for the best Italian M.Sc. thesis in artificial intelligence.

He is the **research manager** (responsabile) **of the research fellowship** (assegno di ricerca) "Development of adaptive reinforcement learning algorithms for control of industrial plants" at DEIB Politecnico di Milano.

Technology Transfer

Start-ups and Spin-offs

Co-founder of ML Cube S.r.l. *November 2020*
Accredited spin-off of Politecnico di Milano
Innovative start-up providing cutting-edge solutions for machine learning systems and life-cycle management optimization. (link: https://www.mlcube.com)
Nominated by Fortune Italia as one of the **top 20 Italian AI start-ups** in 2022.
Funding: **"Smart & Start Italia"** granted by Ministero dello Sviluppo Economico (MiSE) to innovative start-ups (~323 k€).

Open-Source Tools

ARLO (Automated Reinforcement Learning Optimizer) *July 2021-*
Description: Open-source Python library for Automated Reinforcement Learning (AutoRL), aimed at making RL accessible by non-expert users automatizing the development of a learning pipeline. Project status: First released in April 2022, still in development. Role: Scientific advisor.
(github: https://github.com/arlo-lib/ARLO, doc: https://arlo-lib.github.io/arlo-lib/index.html)

Development of Products

ML Cube Platform *November 2020-*
Description: B2B innovative Machine Learning product for Life-cycle Management real-time Optimization, aimed at preventing model obsolescence and providing automatic diagnosis and retraining of the models in production.
Project status: First release. Role: Scientific advisor.
(link: https://www.mlcube.com/mlcube-platform)

AD Cube *December 2021-*
Description: AI platform for the optimization of multi-channel advertising campaigns. Project status: In production. Role: Scientific advisor. (link: https://adcube.ai/)
Funding: Winner of **ELISE's 2nd Open Call grants** (link: https://www.elise-ai.eu/events/elise-s-2nd-open-call-grants-funding-to-16-smes-to-develop-ai-based-services-or-applications)

Talks, Seminars, and Presentations

Invited Talks and Seminars

Invited Talk *February 2024 (to happen)*
38th AAAI Conference on Artificial Intelligence - New Faculty Highlights *Vancouver, Canada*
GGS: A++. Title: "Recent Advancements in Inverse Reinforcement Learning"

Invited Talk *December 2023*
1st Symposium on Lifelong Explainable Robot Learning (SYMPLER) *Nürnberg, Germany*
Title: "Explaining Human Intentions through Inverse Reinforcement Learning"

Invited Seminar *May 2023*
AI Seminars - Politecnico di Milano *Milan, Italy*
Title: "Inverse reinforcement learning: a theoretical view".

Invited Talk *September 2022*
ELLIS@Milan Artificial Intelligence Workshop - Bocconi University *Milan, Italy*
Title: "Online Learning in Non-Cooperative Configurable Environments".

Invited Seminar *May 2022*
Università del Piemonte Orientale *Alessandria, Italy*
Title: "Stream Learning in Non-Stationary Environments".

Invited Talk *December 2021*
20th International Conference of the Italian Association for Artificial Intelligence (AIxIA 2021) *Online*
Title: "Exploiting Environment Configurability in Reinforcement Learning".

Invited Seminar *January 2021*
ML Modena Meetup *Online*
Title: "From MAB to RL... and beyond!" (with Francesco Trovò).

Invited Talk *November 2018*
17th International Conference of the Italian Association for Artificial Intelligence (AIxIA 2018) *Trento, Italy*
Title: "Compatible Reward Inverse Reinforcement Learning".

Seminar *November 2017*
Politecnico di Milano - Dipartimento di Elettronica, Informazione e Bioingegneria *Milan, Italy*
Title: "Distributional Reinforcement Learning".

International Conference and Workshop Talks

Conference Talk *July 2023*
The Fortieth International Conference on Machine Learning (ICML 2023) *Honolulu, HI, US*
GGS: A++. Title: "Towards Theoretical Understanding of Inverse Reinforcement Learning".

Conference Talk *April 2023*
The 26th International Conference on Artificial Intelligence and Statistics (AISTATS 2023) *Valencia, Spain*
GGS: A+. Title: "A Tale of Sampling and Estimation in Discounted Reinforcement Learning".

Conference Talk *July 2022*
Thirty-ninth International Conference on Machine Learning (ICML 2022) *Baltimore, MD, US*
GGS: A++. Title: "Stochastic Rising Bandits".

Spotlight Conference Talk   *December 2021*
Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS 2021)   *Online*
GGS: A++. Title: "Subgaussian and Differentiable Importance Sampling for Off-Policy Evaluation and Learning".

Workshop Talk   *December 2021*
Deep Reinforcement Learning Workshop - NeurIPS 2021   *Online*
Title: "Policy Optimization via Optimal Policy Evaluation".

Conference Talk   *September 2021*
European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in
Databases (ECML-PKDD 2021)   *Online*
GGS: A. Title: "Policy Space Identification in Configurable Environments".

Conference Talk   *July 2021*
Thirty-eighth International Conference on Machine Learning (ICML 2021)   *Online*
GGS: A++. Title: "Provably Efficient Learning of Transferable Rewards".

Conference Talk   *February 2021*
35th AAAI Conference on Artificial Intelligence (AAAI 2021)   *Online*
GGS: A++. Title: "Policy Optimization as Online Learning with Mediator Feedback".

Conference Talk   *July 2020*
Thirty-seventh International Conference on Machine Learning (ICML 2020)   *Online*
GGS: A++. Title: "Control Frequency Adaptation via Action Persistence in Batch Reinforcement Learning".

Short Conference Talk   *June 2019*
Thirty-sixth International Conference on Machine Learning (ICML 2019)   *Long Beach, CA, US*
GGS: A++. Title: "Reinforcement Learning in Configurable Continuous Environments".

Long Conference Talk   *July 2018*
Thirty-fifth International Conference on Machine Learning (ICML 2018)   *Stockholm, Sweden*
GGS: A++. Title: "Configurable Markov Decision Processes".

Poster Presentation at International Conferences

He presented his contributions in several international conferences: ICML 2023 (3 posters - Honolulu, HI, US), AISTATS 2023 (1 poster - Valencia, Spain), ICML 2022 (2 posters - Baltimore, MD, US), NeurIPS 2021 (2 posters - Online), ICML 2021 (1 poster - Online), AAAI 2021 (1 poster - Online), NeurIPS 2019 (1 poster - Vancouver, Canada), ICML 2019 (2 posters - Long Beach, CA, US), NeurIPS 2018 (1 poster - Montreal, Canada), ICML 2018 (1 poster - Stockholm, Sweden), NIPS 2017 (1 poster - Long Beach, CA, US).

Activities and Services[6]

Referee Activities

Registered to REPRISE   *October 2023-*
Register of Expert Peer Reviewers for Italian Scientific Evaluation - MUR

Editorial Activities

Associate Editor   *April 2023-*
IEEE Robotics and Automation Letters (RA-L) - SJR: Q1

Organization Activities

Program Chair and Member of the Organizing Committee   *September 2022*
15th European Workshop on Reinforcement Learning (EWRL 2022)   *Milan, Italy*
Responsibilities: Reviewer recruitment, call for papers writing, supervision of the review and decision process, definition of the meeting schedule, chairing a session.
(link: https://ewrl.wordpress.com/ewrl15-2022/)

Session Chair   *September 2021*

---

[6]For both journal and conference publications, the rating is the most recent available. For Scimago Ranks (SJR), the "Artificial Intelligence" quartile is reported if no further specification.

European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases 2021 (ECML-PKDD) - GGS: A   *Online*
Responsibilities: Chairing a Reinforcement Learning session (4 paper presentations).

(SENIOR) PROGRAM COMMITTEE MEMBER AND REVIEWER FOR INTERNATIONAL CONFERENCES

Senior Program Committee member   *2021*
International Joint Conference on Artificial Intelligence (IJCAI) - GGS: A++

Senior Program Committee member   *2022*
AAAI Conference on Artificial Intelligence (AAAI) - GGS: A++

Program Committee member   *2020, 2022*
International Joint Conference on Artificial Intelligence (IJCAI) - GGS: A++

Reviewer   *2019-2023*
International Conference on Machine Learning (ICML) - GGS: A++

Reviewer   *2019-2023*
Neural Information Processing Systems (NeurIPS) - GGS: A++

Program Committee member   *2020-2021,2023*
European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD) - GGS: A

Program Committee member   *2020-2021,2024*
AAAI Conference on Artificial Intelligence (AAAI) - GGS: A++

Reviewer   *2020-2022*
International Conference on Artificial Intelligence and Statistics (AISTATS) - GGS: A+

Reviewer   *2021-2024*
International Conference on Learning Representations (ICLR) - GGS: A++

REVIEWER FOR INTERNATIONAL JOURNALS

Reviewer   *2023*
Scientific Report (Nature) - SJR: Q1 (Multidisciplinary)

Reviewer   *2023*
PLOS ONE - SJR: Q1 (Multidisciplinary)

Reviewer   *2023*
Expert Systems with Applications (ESWA) - SJR: Q1

Reviewer   *2023*
IEEE Transactions on Artificial Intelligence (TAI) - SJR: Q1

Reviewer   *2023*
IEEE Robotics and Automation Letters (RA-L) - SJR: Q1

Reviewer   *2023*
IEEE Transactions on Neural Networks and Learning Systems (TNNLS) - SJR: Q1

Reviewer   *2022*
Journal of Machine Learning Research (JMLR) - SJR: Q1

Reviewer   *2022-2023*
Transactions on Machine Learning Research (TMLR)

Reviewer   *2022*
Engineering Applications of Artificial Intelligence (EAAI) - SJR: Q1

Reviewer   *2022*
Entropy (MDPI) - SJR: Q2 (Information Systems)

Reviewer   *2021*
Frontiers in Artificial Intelligence - SJR: Q2

Reviewer   *2021*

Intelligenza Artificiale - SJR: Q3

Reviewer   *2021*
IEEE Transactions on Intelligent Transportation Systems (ITS) - SJR: Q1 (Computer Science Applications)

Reviewer   *2019, 2022*
Journal of Artificial Intelligence Research (JAIR) - SJR: Q2

Reviewer   *2019*
Machine Learning (Springer) - SJR: Q1

Reviewer   *2018*
IEEE Transactions on Cognitive and Developmental Systems (TCDS) - SJR: Q1

MEMBERSHIPS OF SCIENTIFIC SOCIETIES

ELLIS Society Member   *2021-*
European Laboratory for Learning and Intelligent Systems

CLAIRE Supporter   *2023-*
Confederation of Laboratories for Artificial Intelligence Research in Europe

AAAI Member   *2020-*
Association for the Advancement of Artificial Intelligence

AIxIA Member   *2018-2022*
Associazione Italiana per l'Intelligenza Artificiale

LIST OF PUBLICATIONS[7]

BOOKS

[B1] <u>Alberto Maria Metelli</u>. *Exploiting environment configurability in reinforcement learning*, volume 361 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2022.
(link: https://doi.org/10.3233/FAIA361).

BOOK CHAPTERS

[H1] <u>Alberto Maria Metelli</u>. "Configurable Environments in Reinforcement Learning: An Overview". In **Special Topics in Information Technology**, pages 101–113, Cham, 2022. Springer International Publishing.
(link: https://doi.org/10.1007/978-3-030-85918-3_9).

REFEREED INTERNATIONAL JOURNAL ARTICLES

[J1] Riccardo Poiani, Ciprian Stirbu, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Optimizing Empty Container Repositioning and Fleet Deployment via Configurable Semi-POMDPs". **IEEE Transactions on Intelligent Transportation Systems**, 2023. SJR 2022: **Q1** (Computer Science Applications).
(link: https://doi.org/10.1109/TITS.2023.3329677).

[J2] Gianluca Drappo, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "An Option-Dependent Analysis of Regret Minimization Algorithms in Finite-Horizon Semi-MDP". **Transactions on Machine Learning Research**, 2023.
(link: https://openreview.net/forum?id=VP9p4u9jAo).

[J3] Filippo Fedeli, <u>Alberto Maria Metelli</u>, Francesco Trovò, and Marcello Restelli. "IWDA: Importance Weighting for Drift Adaptation in Streaming Supervised Learning Problems". **IEEE Transactions on Neural Networks and Learning Systems - Special Issue on Stream Learning**, 34(10):6813–6823, 2023. CORE 2020: **A***. SJR 2022: **Q1**.
(link: https://doi.org/10.1109/TNNLS.2023.3265524).

[J4] Marco Mussi, Davide Lombarda, <u>Alberto Maria Metelli</u>, Francesco Trovó, and Marcello Restelli. "ARLO: A framework for Automated Reinforcement Learning". **Expert Systems with Applications**, 224:119883, 2023. CORE 2020: B. SJR 2022: **Q1**.
(link: https://doi.org/10.1016/j.eswa.2023.119883).

[J5] <u>Alberto Maria Metelli</u>. "A Unified View of Configurable Markov Decision Processes: Solution Concepts, Value Functions, and Operators". **Intelligenza Artificiale**, 16(2):165–184, 2022. SJR 2022: Q3. (*Invited publication as winner of the "Premio Neodottori di Ricerca Marco Cadoli 2021"*).
(link: https://doi.org/10.3233/IA-220140).

[J6] <u>Alberto Maria Metelli</u>, Guglielmo Manneschi, and Marcello Restelli. "Policy space identification in configurable environments". **Machine Learning**, 111(6):2093–2145, 2022. CORE 2020: **A**. SJR 2022: **Q1**.
(link: https://doi.org/10.1007/s10994-021-06033-3).

[J7] <u>Alberto Maria Metelli</u>, Matteo Pirotta, Daniele Calandriello, and Marcello Restelli. "Safe Policy Iteration: A Monotonically Improving Approximate Policy Iteration Approach". **Journal of Machine Learning Research**, 22(97):1–83, 2021. CORE 2020: **A***. SJR 2021: **Q1**.
(link: http://jmlr.org/papers/v22/19-707.html).

[J8] Amarildo Likmeta, <u>Alberto Maria Metelli</u>, Giorgia Ramponi, Andrea Tirinzoni, Matteo Giuliani, and Marcello Restelli. "Dealing with multiple experts and non-stationarity in inverse reinforcement learning: an application to real-life problems". **Machine Learning**, 110(9):2541–2576, 2021. CORE 2020: **A**. SJR 2021: **Q1**.
(link: https://doi.org/10.1007/s10994-020-05939-8).

[J9] <u>Alberto Maria Metelli</u>, Matteo Papini, Nico Montali, and Marcello Restelli. "Importance Sampling Techniques for Policy Optimization". **Journal of Machine Learning Research**, 21(141):1–75, 2020. CORE 2020: **A***. SJR 2020: **Q1**.
(link: http://jmlr.org/papers/v21/20-124.html).

[J10] Amarildo Likmeta, <u>Alberto Maria Metelli</u>, Andrea Tirinzoni, Riccardo Giol, Marcello Restelli, and Danilo Romano. "Combining reinforcement learning with rule-based controllers for transparent and general decision-making in autonomous driving". **Robotics and Autonomous Systems**, 131:103568, 2020. CORE 2020: B. SJR 2020: **Q1** (Computer Science Applications).
(link: https://doi.org/10.1016/j.robot.2020.103568).

---

[7]The first author represents the main contributor; in the case of multiple main contributors (equal contribution), they are marked with *. For both journal and conference publications, the rating is the most recent available in the year of publication. For Scimago Ranks (SJR), the "Artificial Intelligence" quartile is reported if no further specification.

[J11] <u>Alberto Maria Metelli</u>, Matteo Pirotta, and Marcello Restelli. "On the use of the policy gradient and Hessian in inverse reinforcement learning". **Intelligenza Artificiale**, 14(1):117–150, 2020. SJR 2020: Q3. (*Invited publication as winner of the "Premio NeoLaureati Leonardo Lesmo 2018"*).
(link: https://doi.org/10.3233/IA-180011).

Refereed International Conferences Papers

[C1] Riccardo Zamboni, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Distributional Policy Evaluation: a Maximum Entropy approach to Representation Learning". In **Advances in Neural Information Processing Systems 36 (NeurIPS)**. 2023. **Acceptance rate: 26.1%**. CORE 2023: **A***. GGS 2021: **A++**.
(link: https://openreview.net/forum?id=o91in9tDEs).

[C2] Riccardo Poiani, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Truncating Trajectories in Monte Carlo Policy Evaluation: an Adaptive Approach". In **Advances in Neural Information Processing Systems 36 (NeurIPS)**. 2023. **Acceptance rate: 26.1%**. CORE 2023: **A***. GGS 2021: **A++**.
(link: https://openreview.net/forum?id=PkKpTK7hJ6).

[C3] <u>Alberto Maria Metelli</u>, Samuele Meta, and Marcello Restelli. "On the Relation between Policy Improvement and Off-Policy Minimum-Variance Policy Evaluation". In **Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI)**, volume 216, pages 1423–1433. PMLR, 2023. **Acceptance rate: 243/778 (31.2%)**. CORE 2023: **A**. GGS 2021: **A**.
(link: https://proceedings.mlr.press/v216/metelli23a.html).

[C4] <u>Alberto Maria Metelli</u>, Filippo Lazzati, and Marcello Restelli. "Towards Theoretical Understanding of Inverse Reinforcement Learning". In **Proceedings of the 40th International Conference on Machine Learning (ICML)**, volume 202, pages 24555–24591. PMLR, 2023. **Acceptance rate: 1827/6538 (27.9%), Oral: 156/6538 (2.39%)**. CORE 2023: **A***. GGS 2021: **A++**.
(link: https://proceedings.mlr.press/v202/metelli23a.html).

[C5] Riccardo Poiani, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Truncating Trajectories in Monte Carlo Reinforcement Learning". In **Proceedings of the 40th International Conference on Machine Learning (ICML)**, volume 202, pages 27994–28042. PMLR, 2023. **Acceptance rate: 1827/6538 (27.9%)**. CORE 2023: **A***. GGS 2021: **A++**.
(link: https://proceedings.mlr.press/v202/poiani23a.html).

[C6] Marco Mussi, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Dynamical Linear Bandits". In **Proceedings of the 40th International Conference on Machine Learning (ICML)**, volume 202, pages 25563–25587. PMLR, 2023. **Acceptance rate: 1827/6538 (27.9%)**. CORE 2023: **A***. GGS 2021: **A++**.
(link: https://proceedings.mlr.press/v202/mussi23a.html).

[C7] <u>Alberto Maria Metelli</u>, Mirco Mutti, and Marcello Restelli. "A Tale of Sampling and Estimation in Discounted Reinforcement Learning". In **Proceedings of The 26th International Conference on Artificial Intelligence and Statistics**, volume 206, pages 4575–4601. PMLR, 25–27 Apr 2023. **Acceptance rate: 496/1689 (29.3%), Notable paper (oral presentation): 32/1689 (1.9%)**. CORE 2023: **A**. GGS 2021: **A+**.
(link: https://proceedings.mlr.press/v206/metelli23a.html).

[C8] Eldowa Khaled Mazen Mahmoud Elsayed, Nicolò Cesa-Bianchi, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Bandits with Stochastic Experts: Towards Instance-Based Optimality". In **2023 IEEE Information Theory Workshop (ITW)**, pages 30–35. 2023. CORE 2023: **B**. GGS 2021: **B**.

[C9] Davide Maran, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Tight Performance Guarantees of Imitator Policies with Continuous Actions". In **The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)**, pages 9073–9080. AAAI Press, 2023. **Acceptance rate: 19.6%**. CORE 2023: **A***. GGS 2021: **A++**.
(link: https://doi.org/10.1609/aaai.v37i8.26089).

[C10] Amarildo Likmeta, Matteo Sacco, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Wasserstein Actor-Critic: Directed Exploration via Optimism for Continuous-Actions Control". In **The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)**, pages 8782–8790. AAAI Press, 2023. **Acceptance rate: 19.6%**. CORE 2023: **A***. GGS 2021: **A++**.
(link: https://doi.org/10.1609/aaai.v37i7.26056).

[C11] Luca Sabbioni, Luca Al Daire, Lorenzo Bisi, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Simultaneously Updating All Persistence Values in Reinforcement Learning". In **The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI)**, pages 9668–9676. AAAI Press, 2023. **Acceptance rate: 19.6%**. CORE 2023: **A***. GGS 2021: **A++**.
(link: https://doi.org/10.1609/aaai.v37i8.26156).

[C12] Riccardo Poiani, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Multi-Fidelity Best-Arm Identification". In **Advances in Neural Information Processing Systems 35 (NeurIPS)**, volume 35, pages 17857–17870. 2022. **Acceptance rate: 2665/10411 (25.6%)**. CORE 2021: **A***. GGS 2021: **A++**.

(link: http://papers.nips.cc/paper_files/paper/2022/hash/71c31ebf577ffdad5f4a74156daad518-Abstract-Conference.html).

[C13] <u>Alberto Maria Metelli</u>, Matteo Pirola, Francesco Trovò, and Marcello Restelli. "Stochastic Rising Bandits". In **Proceedings of the 39th International Conference on Machine Learning (ICML)**, volume 162, pages 15421–15457. PMLR, 2022. **Acceptance rate: 1235/5630 (21.9%)**. CORE 2021: **A***. GGS 2021: **A++**.
(link: https://proceedings.mlr.press/v162/metelli22a.html).

[C14] Giorgio Manganini, Angelo Damiani, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Balancing Sample Efficiency and Suboptimality in Inverse Reinforcement Learning". In **Proceedings of the 39th International Conference on Machine Learning (ICML)**, volume 162, pages 4618–4629. PMLR, 2022. **Acceptance rate: 1235/5630 (21.9%)**. CORE 2021: **A***. GGS 2021: **A++**.
(link: https://proceedings.mlr.press/v162/damiani22a.html).

[C15] Julen Cestero, Marco Quartulli, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Storehouse: a Reinforcement Learning Environment for Optimizing Warehouse Management". In **2022 IEEE World Congress on Computational Intelligence - International Joint Conference on Neural Networks (IJCNN)**, pages 1–9. 2022. CORE 2021: B. GGS 2021: A-.
(link: https://doi.org/10.1109/IJCNN55064.2022.9891985).

[C16] Pierre Liotet, Francesco Vidaich, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Lifelong Hyper-Policy Optimization with Multiple Importance Sampling Regularization". In **The Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI)**, pages 7525–7533. AAAI Press, 2022. **Acceptance rate: 1349/9020 (15.0%)**. CORE 2021: **A***. GGS 2021: **A++**.
(link: https://doi.org/10.1609/aaai.v36i7.20717).

[C17] <u>Alberto Maria Metelli</u>, Alessio Russo, and Marcello Restelli. "Subgaussian and Differentiable Importance Sampling for Off-Policy Evaluation and Learning". In **Advances in Neural Information Processing Systems 34 (NeurIPS)**, pages 8119–8132. 2021. **Acceptance rate: 2344/9122 (25.7%), Spotlight: 260/9122 (2.9%)**. CORE 2021: **A***. GGS 2021: **A++**.
(link: https://proceedings.neurips.cc/paper/2021/hash/4476b929e30dd0c4e8bdbcc82c6ba23a-Abstract.html).

[C18] Giorgia Ramponi, <u>Alberto Maria Metelli</u>, Alessandro Concetti, and Marcello Restelli. "Learning in Non-Cooperative Configurable Markov Decision Processes". In **Advances in Neural Information Processing Systems 34 (NeurIPS)**, pages 22808–22821. 2021. **Acceptance rate: 2344/9122 (25.7%)**. CORE 2021: **A***. GGS 2021: **A++**.
(link: https://proceedings.neurips.cc/paper/2021/hash/c0f52c6624ae1359e105c8a5d8cd956a-Abstract.html).

[C19] <u>Alberto Maria Metelli</u>*, Giorgia Ramponi*, Alessandro Concetti, and Marcello Restelli. "Provably Efficient Learning of Transferable Rewards". In **Proceedings of the 38th International Conference on Machine Learning (ICML)**, volume 139, pages 7665–7676. PMLR, 2021. **Acceptance rate: 1184/5513 (21.5%)**. CORE 2021: **A***. GGS 2021: **A++**.
(link: http://proceedings.mlr.press/v139/metelli21a.html).

[C20] <u>Alberto Maria Metelli</u>*, Matteo Papini*, Pierluca D'Oro, and Marcello Restelli. "Policy Optimization as Online Learning with Mediator Feedback". In **The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI)**, pages 8958–8966. AAAI Press, 2021. **Acceptance rate: 1692/7911 (21.4%)**. CORE 2021: **A***. GGS 2021: **A++**.
(link: https://ojs.aaai.org/index.php/AAAI/article/view/17083).

[C21] <u>Alberto Maria Metelli</u>, Flavio Mazzolini, Lorenzo Bisi, Luca Sabbioni, and Marcello Restelli. "Control Frequency Adaptation via Action Persistence in Batch Reinforcement Learning". In **Proceedings of the 37th International Conference on Machine Learning (ICML)**, volume 119, pages 6862–6873. PMLR, 2020. **Acceptance rate: 1088/4990 (21.8%)**. CORE 2020: **A***. GGS 2018: **A++**.
(link: http://proceedings.mlr.press/v119/metelli20a.html).

[C22] Giorgia Ramponi, Amarildo Likmeta, <u>Alberto Maria Metelli</u>, Andrea Tirinzoni, and Marcello Restelli. "Truly Batch Model-Free Inverse Reinforcement Learning about Multiple Intentions". In **Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)**, volume 108, pages 2359–2369. PMLR, 2020. CORE 2020: **A**. GGS 2018: **A+**.
(link: http://proceedings.mlr.press/v108/ramponi20a.html).

[C23] Pierluca D'Oro*, <u>Alberto Maria Metelli</u>*, Andrea Tirinzoni, Matteo Papini, and Marcello Restelli. "Gradient-Aware Model-Based Policy Search". In **The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)**, pages 3801–3808, 2020. **Acceptance rate: 1591/7737 (20.6%)**. CORE 2020: **A***. GGS 2018: **A++**.
(link: https://doi.org/10.1609/aaai.v34i04.5791).

[C24] <u>Alberto Maria Metelli</u>*, Amarildo Likmeta*, and Marcello Restelli. "Propagating Uncertainty in Reinforcement Learning via Wasserstein Barycenters". In **Advances in Neural Information Processing Systems 32 (NeurIPS)**, pages 4335–4347, 2019. **Acceptance rate: 428/6743 (21.2%)**. CORE 2018: **A***. GGS 2018:

A++.
(link: https://papers.nips.cc/paper/8685-propagating-uncertainty-in-reinforcement-learning-via-wasserstein-barycenters).

[C25] Mario Beraha, <u>Alberto Maria Metelli</u>, Matteo Papini, Andrea Tirinzoni, and Marcello Restelli. "Feature Selection via Mutual Information: New Theoretical Insights". In **International Joint Conference on Neural Networks (IJCNN)**, pages 1–9. IEEE, 2019. CORE 2018: A. GGS 2018: B.
(link: https://doi.org/10.1109/IJCNN.2019.8852410).

[C26] <u>Alberto Maria Metelli</u>, Emanuele Ghelfi, and Marcello Restelli. "Reinforcement Learning in Configurable Continuous Environments". In **Proceedings of the 36th International Conference on Machine Learning (ICML)**, volume 97, pages 4546–4555. PMLR, 2019. **Acceptance rate: 773/3424 (22.6%)**. CORE 2018: **A***. GGS 2018: **A++**.
(link: http://proceedings.mlr.press/v97/metelli19a.html).

[C27] Matteo Papini, <u>Alberto Maria Metelli</u>, Lorenzo Lupo, and Marcello Restelli. "Optimistic Policy Optimization via Multiple Importance Sampling". In **Proceedings of the 36th International Conference on Machine Learning (ICML)**, volume 97, pages 4989–4999. PMLR, 2019. **Acceptance rate: 773/3424 (22.6%)**. CORE 2018: **A***. GGS 2018: **A++**.
(link: http://proceedings.mlr.press/v97/papini19a.html).

[C28] <u>Alberto Maria Metelli</u>, Matteo Papini, Francesco Faccio, and Marcello Restelli. "Policy Optimization via Importance Sampling". In **Advances in Neural Information Processing Systems 31 (NeurIPS)**, pages 5447–5459, 2018. **Acceptance rate: 1011/4856 (20.8%), Oral: 30/4856 (0.62%)**. CORE 2018: **A***. GGS 2018: **A++**.
(link: http://papers.nips.cc/paper/7789-policy-optimization-via-importance-sampling).

[C29] <u>Alberto Maria Metelli</u>*, Mirco Mutti*, and Marcello Restelli. "Configurable Markov Decision Processes". In **Proceedings of the 35th International Conference on Machine Learning (ICML)**, volume 80, pages 3488–3497, 2018. **Acceptance rate: 618/2473 (25.0%)**. CORE 2018: **A***. GGS 2018: **A++**.
(link: http://proceedings.mlr.press/v80/metelli18a.html).

[C30] <u>Alberto Maria Metelli</u>, Matteo Pirotta, and Marcello Restelli. "Compatible Reward Inverse Reinforcement Learning". In **Advances in Neural Information Processing Systems 30 (NIPS)**, pages 2047–2056, 2017. **Acceptance rate: 678/3240 (20.9%)**. CORE 2017: **A***. GGS 2017: **A++**.
(link: http://papers.nips.cc/paper/6800-compatible-reward-inverse-reinforcement-learning).

Refereed International Workshop Papers

[W1] Riccardo Poiani, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Pure Exploration under Mediators' Feedback". In **NeurIPS 2023 Workshop on Adaptive Experimental Design and Active Learning in the Real World**, 2023.
(link: https://openreview.net/forum?id=wtDzsitgO8).

[W2] Gianmarco Genalti, Lupo Marsigli, Nicola Gatti, and <u>Alberto Maria Metelli</u>. "Towards Fully Adaptive Regret Minimization in Heavy-Tailed Bandits". In **NeurIPS 2023 Workshop Heavy Tails in Machine Learning**, 2023.
(link: https://openreview.net/forum?id=I00Z75alN6).

[W3] Alessio Russo, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Switching Latent Bandits". In **Sixteenth European Workshop on Reinforcement Learning**, 2023.
(link: https://openreview.net/forum?id=H1Om2g-o8Y).

[W4] Filippo Lazzati, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "On the Sample Complexity of Inverse Reinforcement Learning". In **Sixteenth European Workshop on Reinforcement Learning**, 2023.
(link: https://openreview.net/forum?id=xhOOquCNAZ).

[W5] Alessandro Montenegro, Marco Mussi, Francesco Trovò, Marcello Restelli, and <u>Alberto Maria Metelli</u>. "Stochastic Rising Bandits: A Best Arm Identification Approach". In **Sixteenth European Workshop on Reinforcement Learning**, 2023.
(link: https://openreview.net/forum?id=Ctq0d9LEuT).

[W6] Riccardo Zamboni, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Distributional Policy Evaluation: a Maximum Entropy approach to Representation Learning". In **Sixteenth European Workshop on Reinforcement Learning**, 2023.
(link: https://openreview.net/forum?id=WlMKYT-IYLh).

[W7] Francesco Bacchiocchi, Gianmarco Genalti, Davide Maran, Marco Mussi, Marcello Restelli, Nicola Gatti, and <u>Alberto Maria Metelli</u>. "Online Learning in Autoregressive Dynamics". In **Sixteenth European Workshop on Reinforcement Learning**, 2023.
(link: https://openreview.net/forum?id=YrHpQWpwsy).

[W8]   Gianluca Drappo, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "A Provably Efficient Option-Based Algorithm for both High-Level and Low-Level Learning". In **Sixteenth European Workshop on Reinforcement Learning**, 2023.
(link: https://openreview.net/forum?id=awEqv6vAIC).

[W9]   Riccardo Poiani, Nicole Nobili, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Truncating Trajectories in Monte Carlo Policy Evaluation: an Adaptive Approach". In **Sixteenth European Workshop on Reinforcement Learning**, 2023.
(link: https://openreview.net/forum?id=4LSUKNUlQ7).

[W10] Théo Vincent, <u>Alberto Maria Metelli</u>, Jan Peters, Marcello Restelli, and Carlo D'Eramo. "Parameterized projected Bellman operator". In **ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems**, 2023.
(link: https://openreview.net/forum?id=UnNdjopNeW).

[W11] Alessandro Montenegro, Marco Mussi, Francesco Trovò, Marcello Restelli, and <u>Alberto Maria Metelli</u>. "A Best Arm Identification Approach for Stochastic Rising Bandits". In **ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems**, 2023.
(link: https://openreview.net/forum?id=k6aftfkuad).

[W12] Marco Mussi, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Dynamical Linear Bandits for Long-Lasting Vanishing Rewards". **Complex Feedback in Online Learning Workshop @ ICML 2022**, 2022.

[W13] <u>Alberto Maria Metelli</u>, Matteo Pirola, Francesco Trovò, and Marcello Restelli. "Stochastic Rising Bandits for Online Model Selection". **Complex Feedback in Online Learning Workshop @ ICML 2022**, 2022.

[W14] Amarildo Likmeta, Matteo Sacco, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "Directed Exploration via Uncertainty-Aware Critics". **Decision Awareness in Reinforcement Learning Workshop @ ICML 2022**, 2022.
(link: https://openreview.net/forum?id=u5KTFj-hucI).

[W15] <u>Alberto Maria Metelli</u>, Samuele Meta, and Marcello Restelli. "Policy Optimization via Optimal Policy Evaluation". **Deep Reinforcement Learning Workshop - NeurIPS 2021**, 2021.
(link: https://openreview.net/forum?id=iseuu3iLqn).

[W16] <u>Alberto Maria Metelli</u>, Alessio Russo, and Marcello Restelli. "Subgaussian Importance Sampling for Off-Policy Evaluation and Learning". **ICML-21 Workshop on Reinforcement Learning Theory**, 2021.
(link: https://lyang36.github.io/icml2021_rltheory/camera_ready/7.pdf).

[W17] Giorgia Ramponi*, <u>Alberto Maria Metelli*</u>, and Marcello Restelli. "Efficient Inverse Reinforcement Learning of Transferable Rewards". **ICML-21 Workshop on Reinforcement Learning Theory**, 2021.
(link: https://lyang36.github.io/icml2021_rltheory/camera_ready/22.pdf).

[W18] Giorgia Ramponi, <u>Alberto Maria Metelli</u>, Alessandro Concetti, and Marcello Restelli. "Online Learning in Non-Cooperative Configurable Markov Decision Process". **AAAI-21 Workshop on Reinforcement Learning in Games**, 2021.
(link: http://aaai-rlg.mlanctot.info/papers/AAAI21-RLG_paper_7.pdf).

[W19] Amarildo Likmeta, <u>Alberto Maria Metelli</u>, Giorgia Ramponi, Andrea Tirinzoni, Matteo Giuliani, and Marcello Restelli. "Handling Non-Stationary Experts in Inverse Reinforcement Learning: A Water System Control Case Study". **Challenges of Real-World RL Workshop @ NeurIPS 2020**, 2020.
(link: https://drive.google.com/file/d/1v3CiRlWtOVJZry15DQdxzeh98UoNAWbA/view).

[W20] Amarildo Likmeta, <u>Alberto Maria Metelli</u>, Andrea Tirinzoni, Riccardo Giol, Marcello Restelli, Danilo Romano, and Andrea Alessandretti. "Autonomous Driving with Reinforcement Learning and Rule-based Policies". **Workshop on AI for Autonomous Driving (AIAD) @ICML 2020**, 2020.
(link: https://drive.google.com/file/d/1ASJa-pOgZ_Z78KTVjrTV1kqqtQ5-RP_w/view).

[W21] Pierluca D'Oro*, <u>Alberto Maria Metelli*</u>, Andrea Tirinzoni, Matteo Papini, and Marcello Restelli. "Gradient-Aware Model-based Policy Search". **Workshop on Meta-Learning (MetaLearn 2019) @NeurIPS 2019**, 2019.
(link: http://metalearning.ml/2019/papers/metalearn2019-doro.pdf).

[W22] <u>Alberto Maria Metelli*</u>, Mirco Mutti*, and Marcello Restelli. "Configurable Markov Decision Processes". **European Workshop on Reinforcement Learning 14 (EWRL 14)**, 2018.
(link: https://ewrl.files.wordpress.com/2018/09/ewrl_14_2018_paper_5.pdf).

[W23] Mattia Bianchi, Federico Cesaro, Filippo Ciceri, Mattia Dagrada, Alberto Gasparin, Daniele Grattarola, Ilyas Inajjar, <u>Alberto Maria Metelli</u>, and Leonardo Cella. "Content-Based Approaches for Cold-Start Job Recommendations". In **Proceedings of the Recommender Systems Challenge 2017**, RecSys Challenge '17, pages 6:1–6:5. ACM, 2017.
(link: http://doi.acm.org/10.1145/3124791.3124793).

[A1] Giorgio Manganini, Angelo Damiani, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "A Novel Inverse Reinforcement Learning Formulation for Sample-Aware Forward Learning". **The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)**, 2022.

[A2] Luca Sabbioni, Luca Al Daire, Lorenzo Bisi, <u>Alberto Maria Metelli</u>, and Marcello Restelli. "All-persistence Bellman Update for Reinforcement Learning". **The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM)**, 2022.

[A3] Matteo Giuliani, Paolo Bonetti, <u>Alberto Maria Metelli</u>, Marcello Restelli, and Andrea Castelletti. "Advancing drought monitoring via feature extraction and multi-task learning algorithms". **EGU General Assembly 2022**, 2022.
(link: https://meetingorganizer.copernicus.org/EGU22/EGU22-5130.html).

[A4] Verónica Torralba, Stefano Materia, Carmen Álvarez Castro, Paolo Bonetti, <u>Alberto Maria Metelli</u>, Marcello Restelli, and Silvio Gualdi. "Seasonal forecasts for hydropower: downscaling of precipitation in South American basins". **EGU General Assembly 2022**, 2022.
(link: https://meetingorganizer.copernicus.org/EGU22/EGU22-13206.html).

[A5] Matteo Giuliani, <u>Alberto Maria Metelli</u>, Andrea Castelletti, and Marcello Restelli. "Advancing drought monitoring via feature extraction". **Earth and Space Science Open Archive**, page 1, 2021.
(link: https://www.essoar.org/doi/abs/10.1002/essoar.10508913.1).

Theses

[T1] <u>Alberto Maria Metelli</u>. "Compatible Reward Inverse Reinforcement Learning". Master's thesis, Politecnico di Milano, July 2017.
(link: http://hdl.handle.net/10589/135141).

[T2] <u>Alberto Maria Metelli</u>. "Exploiting Environment Configurability in Reinforcement Learning". PhD thesis, Politecnico di Milano, March 2021.
(link: http://hdl.handle.net/10589/170616).

Autorizzo al trattamento dati ai sensi del GDPR 2016/679 del 27 aprile 2016 (Regolamento Europeo relativo alla protezione delle persone fisiche per quanto riguarda il trattamento dei dati personali).